DREAMING READOUT

Dreaming for Agent Memory

Review what an agent remembers, offline, against what happened

A visual readout of dreaming: an asynchronous job that checks an agent's memory against past session transcripts and writes a reorganized memory store. Status, mechanics, evidence, and architect's notes.

FIRST-PARTY POST AND DOCS / INDEPENDENT LUMINITY SYNTHESIS



Dreaming

An offline job that checks an agent's memory against its past sessions. Research preview, access by request.

01

Can you use it?

✓ Claude Managed Agents	REQUIRED
✓ Access to Dreams, granted on request	REQUIRED
✓ A supported model	REQUIRED
✓ A memory store plus 1 to 100 past sessions	REQUIRED
✓ Beta headers on each call	REQUIRED
✕ Zero data retention (ZDR) and HIPAA	NOT ELIGIBLE

Dreams · preview · Claude API only

Memory stores · beta

02

How a dream runs



INPUT STORE

Never modified, so you can discard the result

BILLING

Standard token rates; scales with the sessions you feed it

STATES

pending, running, completed, failed, canceled

03

Conditions and limits

Steering — OPTIONAL

Instructions guide the synthesis. Not a line editor: for targeted edits, use the Memory Stores API on the output store.

Models · Dreams page list — CONFIRM LIVE

Supported: Fable 5.1, Fable 5, Opus 5.5, Opus 5, Sonnet 5.5, Sonnet 5, Opus 4.8, Opus 4.7, Sonnet 4.6.
Others are not on the list, including models the memory page lists, such as Haiku 5.5 and Opus 4.6 (Luminity's comparison).

If an input disappears — FAILS

Archiving or deleting an input store, or deleting an input session, mid-run makes the dream fail. A failed or canceled run leaves a partial store.

DATA HANDLING

- Session transcripts persist until you delete them
- Not ZDR or HIPAA eligible, per the docs
- Default rate limits apply; confirm status on the live pages

Sources: Claude Platform Docs, "Dreams," platform.claude.com/docs/en/managed-agents/dreams (retrieved Nov. 8, 2026). Claude Platform Docs, "Using agent memory," platform.claude.com/docs/en/managed-agents/memory (retrieved Nov. 8, 2026). Beta headers: dream calls send dreaming-2026-04-21 with managed-agents-2026-04-01; memory-store calls use agent-memory-2026-07-22 alone (400 if combined). ZDR and HIPAA chips copy the docs' data-handling lines. The flow graphic is Luminity's.

Attribution: Anthropic, Claude and Claude Code are trademarks of Anthropic PBC. Independent analysis of public documentation; not sponsored or endorsed by Anthropic. © 2026 Luminity Digital, Inc.

Memory stores

What dreaming reads and writes: a workspace-scoped collection of text documents. In beta.

01

How a store is used

✓ Mounted as a directory under /mnt/memory/	MOUNT
✓ The agent uses its ordinary file tools	FILE TOOLS
✓ Attach at session creation, not to a running session	AT CREATION
✓ read_write (default) or read_only	ACCESS MODE
✓ A content hash can make a stale write fail	SAFE EDITS
✓ Separate stores for separate owners or access rules	UP TO 8

Access mode is enforced at the filesystem level.

02

Limits

100 kB	10,000	8	30 days
per memory	memories per store	stores per session	versions kept

AUDIT TRAIL

Every change makes an immutable version; recent versions of a live memory are always kept

REDACT

Scrubs a version’s content, keeps who did what and when. The current version can’t be redacted

LONGER HISTORY

Export versions through the API

03

When a store fills up

At 10,000 memories — FAILS

Writes to new memories fail.

Run a dream — DREAM

Consolidate into a separate output store, as the docs suggest.

Switch sessions — ARCHIVE

Point sessions at the new store, then archive or delete the original.

Availability — BETA

Beta on every model. Listed for the Claude API, Claude Platform on AWS, and Claude Models on Google Cloud (Anthropic managed), with some self-hosted limits.

DATA HANDLING • Memory stores: not ZDR eligible | • Not HIPAA eligible | • Per the memory docs’ data-handling line

Sources: Claude Platform Docs, “Using agent memory,” platform.claude.com/docs/en/managed-agents/memory (retrieved Nov. 8, 2026). Sections: “Overview,” “Attach a memory store to a session,” “How the agent accesses memory,” “Audit memory changes,” “Best practices,” and “Availability.”

Attribution: Anthropic, Claude and Claude Code are trademarks of Anthropic PBC. Independent analysis of public documentation; not sponsored or endorsed by Anthropic. © 2026 Luminity Digital, Inc.

Why consolidate offline

Agents write memory in the moment. A dream checks it afterward, with the transcripts as ground truth.

01

Capture, then consolidate

1 **Memory notes** — what the agent wrote in-session

2 **Transcripts** — the full event log: ground truth

3 **Dream** — tests notes against transcripts

4 **Revised notes** — kept or discarded after review

Review — before sessions use them

WHY NOT AT CAPTURE	A WRONG NOTE	THE TRADE-OFF	THE POST SAYS
Writing notes in the moment draws attention from the task	"Can be read, and trusted, by Claude in a later session"	Outputs may not be immediately accessible to Claude, and it requires compute	It may be important to focus dreaming on information likely to change, or to tune its cadence

02

Memory layers

Injected context
small set of notes loaded at the start of every session

Memory representations
structures that help navigate memory, such as linked wiki pages

Memory notes
what Claude judged worth recording in a session

Transcripts
session logs and transcripts, searched when needed

03

What a dream can do

✓	Corrects stale or contradicted entries	DOCS + POST
✓	Consolidates duplicate entries	DOCS + POST
✓	Surfaces patterns that span sessions	DOCS + POST
!	Builds representations, such as linked wiki pages	POST ONLY
!	Promotes key notes to the injected top layer	POST ONLY

SOURCE LIMIT • The Dreams page never mentions wikis, representations, or promotion | • Layer names are Luminity's rendering

Sources: Anthropic, "The components of agent memory," anthropic.com/research/components-of-agent-memory (published Nov. 6, 2026; retrieved Nov. 8, 2026). Claude Platform Docs, "Dreams," platform.claude.com/docs/en/managed-agents/dreams (retrieved Nov. 8, 2026). Sections: introduction, "Memory notes," "Dreaming can correct memory," "Thinking backward, looking forward"; the Dreams page "How it works." Layout and layer labels are Luminity's; the post's figures were read as text. DOCS = Dreams page, POST = research post.

Attribution: Anthropic, Claude and Claude Code are trademarks of Anthropic PBC. Independent analysis of public documentation; not sponsored or endorsed by Anthropic. © 2026 Luminity Digital, Inc.

Evidence: correction

Two examples from Claude playing Pokémon Red with a mounted memory folder. The post presents them as examples, not statistics.

01

A dream catches a wrong note

1
Wrong note
Wrote that the PokéCenter landmark was on the east side (it was not), then "Mt Moon DONE."

2
Wrong lesson
Walked off a one-way ledge, landed back on the route, and concluded it had been on the right side.

3
Dream
Checked the note against the transcript, caught the error, and rewrote it.

With the corrected note
REACHED

Claude went back into the cave and eventually reached Cerulean City.

With the uncorrected note
DID NOT

It dropped back onto that route three times, and the run ended without reaching Cerulean City.

02

What to save · around step 10,000

Sonnet 3.7



413 memory files in 225 folders · 1 badge

Fable 5



9 files in 2 folders, plus global.md · 4 badges

What each kept

Sonnet 3.7 assumed everything might matter, so its record kept going stale. Fable 5 used memory for what the game doesn't show: strategy, plans, and tested rules.

Steering · CL-Bench

General guidance was enough for Opus 4.7 and Opus 5. Specific prompting (read memory first, test theories) helped older models such as Sonnet 4.6 the most. A fixed rule of "record only failures" cost Opus 5 and Opus 4.7 much of their memory gain.

SCOPE

- Single examples: one memory folder per model, one corrected note
- Bars are scaled by Luminity
- CL-Bench and the Mt. Moon correction ran in Managed Agents

Sources: Anthropic, "The components of agent memory," anthropic.com/research/components-of-agent-memory (published Nov. 6, 2026; retrieved Nov. 8, 2026). Sections: "Claude has learned to judge salience," "Steering Claude to write memory" (the steering box comes from the post's CL-Bench test, not Pokémon), "Dreaming can correct memory." Bars are scaled by Luminity, with a minimum width for the smaller count. The flow graphic is Luminity's rendering.

Attribution: Anthropic, Claude and Claude Code are trademarks of Anthropic PBC. Independent analysis of public documentation; not sponsored or endorsed by Anthropic. © 2026 Luminity Digital, Inc.

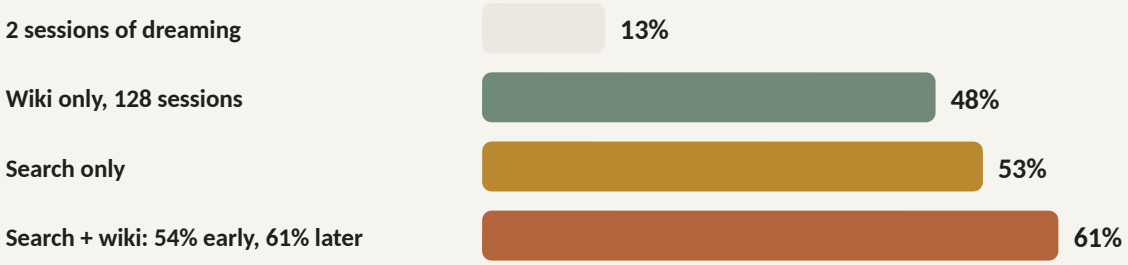
Evidence: wikis

A wiki is a set of linked pages Claude writes about a subject. It helped most when answers span several documents.

01

Score as % of graded criteria

25 QUESTIONS · ~350,000 FILES · ROUGHLY 420M TOKENS



Wiki-only (48%) was still below search alone (53%). Search plus wiki used about 20,000 fewer tokens per question (baseline not stated).

02

Merger questions that span documents

24-QUESTION TEST · SUBSET NEEDING SYNTHESIS, JUDGMENT OR MATH (SIZE NOT STATED)



! Single-document tasks: no clear wiki gain — ONE LEGAL BRIEF

! Wiki grew larger than the documents themselves — ANOTHER TASK

03

More results from the post

Who writes

OWNERSHIP

Several agents writing one shared table kept 41 of 57 graded contract facts. One agent per contract, quoting terms word for word, kept 57 of 57.

Missing pages

SECOND PASS

Sometimes Claude linked to wiki pages that did not exist yet. On one task a second dreaming pass found nine links to one missing page, repointed them to an existing page, and added three new pages, including an index.

Better instructions

LOOP

Claude found them itself by looping: reflect on failures, propose a change, rebuild the wiki. The brief from the improved wiki met 100% of criteria, up from 72%.

SCOPE

- Opus 5 agents built the wiki for the 25-question corpus; the post names no model for the merger and legal-brief tests

READ AS

- Not stated to use the Dreams API | • Small question sets | • A direction, not a forecast

Sources: Anthropic, "The components of agent memory," anthropic.com/research/components-of-agent-memory (published Nov. 6, 2026; retrieved Nov. 8, 2026). Section: "Dreaming can build new memory representations," "Dreaming can focus memory on a specific goal," and "Dreaming can fill in gaps a session leaves behind." Score rises are as stated in the post's text: 13% to 48% (the post gives 48% as the best wiki-only score; that the 13% point is wiki-only is inferred) and Search + wiki 54% to 61%. Bars are scaled by Luminity.

Attribution: Anthropic, Claude and Claude Code are trademarks of Anthropic PBC. Independent analysis of public documentation; not sponsored or endorsed by Anthropic. © 2026 Luminity Digital, Inc.

Architect’s notes

Design rules from the docs and the post, and the related feature that already ships.

01

Before you build

- ✓ **Review the new store before sessions use it** — A dream writes to a separate store — DOCS
- ✓ **Give each subject one writer** — Use read-only mounts and hash preconditions for anything shared — POST + DOCS
- ✓ **Mount reference material read-only** — Untrusted input written to a read-write store becomes trusted memory later — DOCS
- ✓ **Keep injected context small** — Promoted notes cost tokens every session; longer instruction files reduce adherence — POST
- ✓ **Check the supported model list** — Dreams run on select models; the list sits on the Dreams page — DOCS
- ! **Regulated data needs review first** — Not ZDR or HIPAA eligible; transcripts persist until you delete them — DOCS + LUMINITY

02

Related: Claude Code auto memory

On by default · local sessions | **Setting** · autoMemoryEnabled

- ✓ First 200 lines or 25 KB of MEMORY.md load — AT START
- ✓ MEMORY.md is an index; topic files load when needed — ON DEMAND
- ✓ Toggle with /memory — CONTROL
- ✓ Off by default in self-hosted-environment sessions outside Claude Tag — EXCEPTION
- ✓ Machine-local: not shared across machines or cloud environments — LOCAL

The post cites this as the same small-top-layer idea.

NOT COVERED

- Outcomes and multiagent orchestration (Anthropic’s May 19, 2026 launch post for Managed Agents)
- Customer results

LIMITS

- Informational, not legal or compliance advice
- Status and limits as of Nov. 8, 2026; they can change
- The post’s charts were not read

Sources: Claude Platform Docs, platform.claude.com/docs/en/managed-agents/ (dreams, memory); Claude Code Docs, code.claude.com/docs/en/memory; Anthropic, “The components of agent memory.” Retrieved Nov. 8, 2026. Rows are Luminity’s synthesis; the tags show which source each rests on. DOCS = Dreams and memory pages, POST = research post.

Attribution: Anthropic, Claude and Claude Code are trademarks of Anthropic PBC. Independent analysis of public documentation; not sponsored or endorsed by Anthropic. © 2026 Luminity Digital, Inc.